arXiv:1602.07780v4 [q-bio.PE] 28 Nov 2016

DIMENSIONAL REDUCTION FOR THE GENERAL MARKOV MODEL ON PHYLOGENETIC TREES

JEREMY G SUMNER

Abstract. We present a method of dimensional reduction for the general Markov model of sequence evolution on a phylogenetic tree. We show that taking certain linear combinations of the associated random variables (site pattern counts) reduces the dimensionality of the model from exponential in the number of extant taxa, to quadratic in the number of taxa, while retaining the ability to statistically identify phylogenetic divergence events. A key feature is the identification of an invariant subspace which depends only bilinearly on the model parameters, in contrast to the usual multi-linear dependence in the full space. We discuss potential applications including the computation of split (edge) weights on phylogenetic trees from observed sequence data.

1. Introduction

Phylogenetics is the suite of mathematical and computational methods available to biologists for the inference of the evolutionary history of extant species. Typical data inputs to these methods are molecular sequences in various forms, including DNA, sequences of amino acids, codons, and proteins (usually first aligned and then summarised into site pattern counts), and even gene orderings at the scale of whole genomes. Underlying many modern methods (particularly likelihood and Bayesian approaches) is a stochastic model of molecular state evolution; most typically constructed as a Markov process on candidate evolutionary trees. The reader is referred to the texts [13, 24, 31] for excellent backgrounds on the mathematical, statistical, and computational aspects of phylogenetics.

Over the last two decades or so there has been much mathematical progress emphasizing natural *algebraic* structures that underlie phylogenetic models. These range from algebraic geometry arising from model parameterization maps [2, 25], to the combinatorial group theory inherent in genome rearrangement models [16], to methods of discrete Fourier transforms [18]. Concurrently, our research has explored the role of (matrix) Lie groups [27] together with their associated representation theory and invariant functions [21, 26, 29], in particular defining what we refer to as "Markov invariants"[1] (we provide a recent review in [20]).

Date: November 29, 2016.

[1]Not to be confused with "phylogenetic invariants"; the distinction will be discussed below.

In their most abstract setting, Markov invariants provide the means to reduce the high dimensionality of phylogenetic models (determined by the number of observable site patterns) from exponential in the number of taxa to one-dimensional subspaces spanned by individual polynomials (from the statistical point of view, these are functions on the random variables given by site pattern counts in molecular sequence alignments). Unfortunately, due to algebraic combinatorial explosion, it is only feasible to explore these polynomials explicitly for small cases, and this has only been achieved so far for four taxa — the so-called "quartet" case [19, 30].

Even though the mathematical properties of Markov invariants are rather elegant, and in themselves quartets are of fundamental importance to phylogenetics, the practical utility of this approach becomes questionable when biologists are interested in the inference of large phylogenetic trees with many, possibly thousands, of taxa. For these and other reasons, it is perhaps not surprising that the emergence of algebraic approaches has not been accompanied by an upsurge in usage by practicing biologists.

In this paper, we will explore a generalization which compromises on the dimensional reduction provided by Markov invariants slightly: we reduce the dimensionality of the model from exponential in the number of taxa to *quadratic* in the number of taxa. This reduction comes with the significant benefit that the method we describe is generally applicable to any number of taxa using standard computational techniques currently available in most programming languages (linear transformations and singular value decomposition SVD), with no further theoretical work required. The approach hinges on a simple algebraic observation regarding the embedding of Markov matrices into the matrix group of linear invertible affine transformations, and the identification of a particular invariant subspace after distributing over Kronecker (or tensor) products.

The paper proceeds as follows. In §2 we give the required background on the algebraic structures underlying phylogenetic models. Specifically, we discuss the so-called "flattenings" and associated rank properties. In §3 we present our main result of dimensional reduction for phylogenetic models. In §4 we discuss potential strategies for applying the dimensional reduction in the practical setting of phylogenetic tree inference.

2. Background

We assume we have L biological taxa labelled as $1, 2, \ldots, L$ together with homologous (aligned) molecular states drawn from a finite state space κ. For example, we have $|\kappa| = 4$ for DNA and $|\kappa| = 20$ for amino acids. Since our results are independent of the size of the state space, we will generically write $|\kappa| = k$. To produce a phylogenetic model, we fix (i) a (rooted) binary tree $\mathcal{T}$ with L leaves (degree 1 vertices), (ii) an initial probability distribution $(\pi_i)_{i\in\kappa}$, and (iii) Markov

matrices M_e for each edge $e \in \mathcal{T}$. (Throughout this paper, we will follow the convention that Markov matrices have unit-*column* sums, as opposed to unit-row sums.) These inputs then produce a probability distribution of site patterns $P = (p_{i_1 i_2 \ldots i_L})_{i_j \in \kappa}$ at the leaves of the tree and we refer to any such distribution as having "arisen" under the general Markov model on the tree $\mathcal{T}$. This process is illustrated in Figure 1.

FIGURE 1. Graphical illustration of the construction of a site pattern distribution $P = (p_{i_1 i_2 i_3 i_4 i_5 i_6 i_7})_{i_j \in \kappa}$ on a phylogenetic tree $\mathcal{T}$ with seven taxa. Input data is a set of Markov matrices $\{M_e\}_{e \in \mathcal{T}}$ and an initial distribution $(\pi_i)_{i \in \kappa}$.

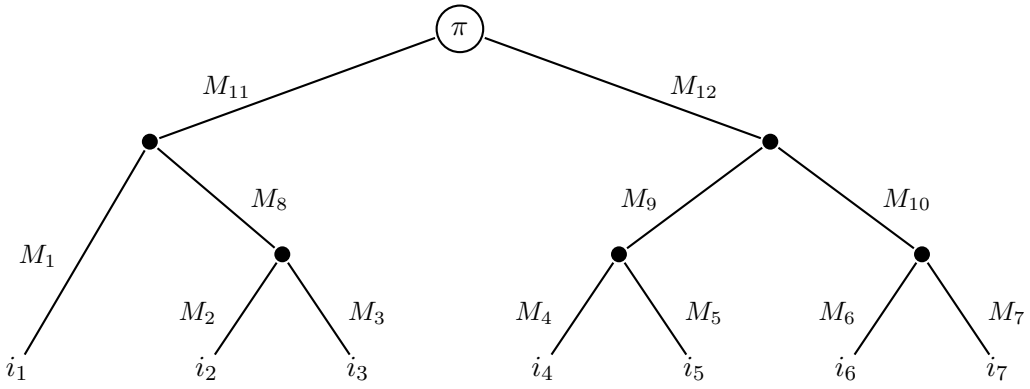

Even though such a model specifies a location for the root of the tree, it is generally the case that the choice of root either does not affect the output distribution in any way [12], or for each alternative root position there is a corresponding alternative choice of model parameters which produce the same distribution. Thus, without loss of generality, we may remove the root vertex from $\mathcal{T}$ and hitherto consider $\mathcal{T}$ to be a binary tree with all vertices of degree 3 or 1. For simplicity we will refer to both the vertices of degree 1 *and* their adjacent edges as "leaves" (context will resolve any ambiguity).

Consider an $m|n$ "split" $A|B$ of the taxon set $\{1, 2, \ldots, L\}$. This is defined as a bipartition $\{A, B\}$ with $A, B \subseteq \{1, 2, \ldots, L\}$, $A \cap B = \emptyset$, $A \cup B = \{1, 2, \ldots, L\}$, $|A| = m$, and $|B| = n$, so $m + n = L$. Suppose we fix an (unrooted) binary tree $\mathcal{T}$ with leaf set $\{1, 2, \ldots, L\}$. Then each edge in $\mathcal{T}$ gives rise to a split $A|B$ via deletion of the edge and the resulting partitioning of the leaf vertices. Correspondingly, for each split $A|B$ we may define a so-called "flattening" of the array $P = (p_{i_1 i_2 \ldots i_L})_{i_j \in \kappa}$, as follows[2]. Supposing $A = \{s_1, s_2, \ldots, s_m\}$ and $B = \{t_1, t_2, \ldots, t_n\}$, the flattening $\mathrm{Flat}_{A|B}(P)$ is the $k^m \times k^n$ matrix with rows indexed by the strings $i_{s_1} i_{s_2} \ldots i_{s_m}$, columns indexed by strings $i_{t_1} i_{t_2} \ldots i_{t_n}$, and entries $p_{i_1 i_2 \ldots i_L}$. Although the row and column indices

[2]Illustrative examples of flattenings are given in the introduction of [2].

may be ordered in any way we choose without affecting the present discussion, it is convenient to take the ordering which is consistent with the tensor/Kronecker products discussed in the next section.

This construction gives rise to the so-called "edge invariants" since it can be shown that *the flattening* $Flat_{A|B}(P)$ *has matrix rank at most* k *if* $A|B$ *is a split in the tree and at least* k^2 *otherwise.* (These facts are established in [2, 11] for "generic" cases — weak conditions on model parameters need to be assumed.) More generally, [11] gives an algorithm to find $r \in \mathbb{N}$ such that the rank of the flattening is at most k^r and, up to this bound, maximal rank is attained in generic cases.

The edge invariants are defined as the $(k+1)\times(k+1)$ minors of the flattening $\mathrm{Flat}_{A|B}(P)$, which are seen to vanish if the edge corresponding to $A|B$ occurs in $\mathcal{T}$. These minors form polynomials in the variables $(p_{i_1 i_2 \ldots i_L})_{i_j \in \kappa}$ and are specific cases of what are known as "phylogenetic invariants" in the literature [8, 23], which are defined as polynomials that vanish on probability distributions arising from certain subsets of phylogenetic trees (see [2, 7, 10] for more recent developments). For our purposes, it is important to distinguish "phylogenetic" from "Markov" invariants; these are formally distinct concepts with instances that sometimes, but not always, coincide.

The flattenings and/or edge invariants may be used as a simple test for the occurrence of the edge corresponding to $A|B$ in the tree $\mathcal{T}$ as follows.

(1) Collect the observed molecular sequences into an alignment (using standard, freely available software) to produce a site pattern count array $F = (f_{i_1 i_2 \ldots i_L})_{i_j \in \kappa}$, where $f_{i_1 i_2 \ldots i_L}$ is the number of sites which have state pattern $i_1 i_2 \ldots i_L$;
(2) Suppose that F is a sample $F \sim \text{multinomial}(P, N)$ where P is some pattern distribution arising under the general Markov model of sequence evolution on an (unknown) tree $\mathcal{T}$ and N is the length of the alignment;
(3) Test for suitability of the edge corresponding to $A|B$ belonging to $\mathcal{T}$ by computing the rank of $\mathrm{Flat}_{A|B}(F)$;
(4) If this rank is "close" to k, conclude that $A|B$ is in the tree $\mathcal{T}$.

Ignoring issues of how to numerically estimate the rank ([11] suggests using the SVD decomposition, and [6] have explored using the minors directly), or how one decides what we mean by "close" in a statistically justified way (this is an open problem generally), as a first step this process provides a mathematically elegant test for the evidence of a specific edge in a phylogenetic tree. Significantly, this test is statistically valid no matter what Markov model of molecular state evolution is presupposed (since it is valid for the "general" Markov model). On the down side, the complexity of the rank estimation and comparison problem is exponential in the number of taxa L

since $\mathrm{Flat}_{A|B}(F)$ is a $k^m \times k^n$ matrix. For a large number of taxa $L=m+n$ and finite sequence length, the flattening matrix will clearly be sparse.

The SVD methodology has implemented by [14] and [9] with impressive results. However, these studies restrict attention to quartets only. Most recently [1] presented an efficient implementation of the flattenings and SVD for arbitrary tree sizes using sparse matrix computation.

3. Dimensional reduction

As has been observed in many places [2, 4, 5, 28], phylogenetic models implemented as a Markov chain on a binary tree can be described algebraically using tensor product spaces and/or Kronecker product operations. Following the discussion given in [26], if we fix a site pattern distribution P obtained from the general Markov model on a phylogenetic tree as described in the previous section, we may define an alternative distribution $\tilde{P}$ specified by setting all Markov matrices on the leaves of the tree equal to the identity. As is shown in [26], using tensor notation we may write

$$P = M_1 \otimes M_2 \otimes \ldots \otimes M_L \cdot \tilde{P}. \tag{3-1}$$

Written in terms of an $m|n$ flattening on the split $A|B$, this equation is expressible using matrix products:

$$\mathrm{Flat}_{A|B}(P) = M_A \mathrm{Flat}_{A|B}(\tilde{P}) M_B^T,$$

where M_A is the $k^m \times k^m$ matrix obtained by taking the Kronecker product of the Markov matrices on the leaves belonging to A, and M_B is the $k^n \times k^n$ matrix obtained by taking the Kronecker product of the Markov matrices on the leaves belonging to B.

Further, given that Markov matrices are multiplicative (that is, multiplication of two Markov matrices gives another Markov matrix), we may generalise (3-1) to be understood as the transformation rule

$$P \to M_1 \otimes M_2 \otimes \ldots M_L \cdot P, \tag{3-2}$$

expressing how a distribution P arising from the general Markov model on a phylogenetic tree changes if additional Markov matrices M_i are applied at the leaves of the tree. Expressed in terms of the flattening, this transformation rule becomes

$$\mathrm{Flat}_{A|B}(P) \to M_A \mathrm{Flat}_{A|B}(P) M_B^T. \tag{3-3}$$

Practical importance of this transformation rule follows from the observation that the uniqueness of the binary tree underlying the phylogenetic model is unaffected by the choice of Markov matrices on the leaves (even under the condition that they are set equal to the identity matrix), and this is highlighted by the fact that the corresponding statement is not true for the internal edges of

the tree. To see this, one may refer to Figure 1 where restricting the matrix M_8 to be equal to the identity would imply, as judged by the distribution P, the split $\{2,3\}|\{1,4,5,6,7\}$ is indistinguishable from the alternative splits $\{1,2\}|\{3,4,5,6,7\}$, $\{1,3\}|\{2,4,5,6,7\}$, or $\{1,2,3\}|\{4,5,6,7\}$. Additionally, it is easily seen that, if the relevant Markov matrices are of full rank, *the rank of the flattening is unaffected by the transformation rule (3-3).* The general philosophy we follow is that the transformation rules (3-2) and (3-3) are foundational to algebraic approaches to identifying splits on phylogenetic trees. This aligns with our previous work on Markov invariants [26] which seeks to find polynomials on the variables $(p_{i_1 i_2 \ldots i_L})_{i_j \in \kappa}$ that are invariant (up to scaling) under the transformation rule (3-2).

In this section we show that we can significantly reduce the dimensionality of the matrices involved in the transformation rule (3-3) by implementing a linear change of coordinates on the flattening $\mathrm{Flat}_{A|B}(P)$ and identifying the "sub-flattening" $\widehat{\mathrm{Flat}}'_{A|B}(P)$ as a certain sub-matrix thereof. Crucially, the dimension of the sub-flattening $\widehat{\mathrm{Flat}}'_{A|B}(P)$ is *quadratic* in the number of taxa L and the sub-flattening retains contrasting rank conditions when the true tree $\mathcal{T}$ does or does not contain the edge corresponding to $A|B$.

Before we do this however, the reader should note that the sub-flattening $\widehat{\mathrm{Flat}}'_{A|B}(P)$ is distinct from the related concept of "thin flattenings" introduced in [7], which are specific to each Markov model in the "equivariant" hierarchy presented in [10]. Although the components of a thin flattening have reduced dimension (relative to the dimension of the full flattening $\mathrm{Flat}_{A|B}(P)$), their dimension remains exponential in the number of taxa. Importantly, in the case of the general Markov model, the thin flattening is equal to the original flattening and hence does not produce a dimensional reduction.

To derive the sub-flattening $\widehat{\mathrm{Flat}}'_{A|B}(P)$, we make the simple observation that, under an appropriate similarity transformation, the condition that Markov matrices have unit column-sums[3] is converted into the condition that a certain row has $k-1$ zeros and a single 1. Concretely, we may choose the similarity transformation S so any $k \times k$ Markov matrix takes the form

$$M' := SMS^{-1} = \begin{pmatrix} T & v \\ 0 & 1 \end{pmatrix}, \tag{3-4}$$

where T is a $(k-1)\times(k-1)$ matrix and v is a $k-1$ column vector. Generically (and hence without loss of relevance to the applied setting), we may assume $\det(M) \neq 0$ which in turn implies $\det(T) \neq 0$. The set of such matrices (3-4) then forms what is known as the *affine group* $\mathrm{Aff}(k-1)$ [3, Chap. 1]. This is the (matrix) Lie group of invertible affine linear transformations on $\mathbb{R}^{k-1}$ given

[3]If one prefers to use unit row-sum Markov matrices, an analogous construction is obtained by taking the transpose in what follows.

by the semi-direct product $\mathrm{GL}(k{-}1) \ltimes \mathbb{R}^{k-1}$. (This correspondence under similarity transformations between Markov matrices and the affine group seems to have first been made in [22].)

In applications it is always the case that the entries of Markov matrices M are non-negative — this after all ensures probability conservation of the Markov process. However, we need not worry about these additional conditions to obtain the results we wish to derive here — the above embedding of the transformed matrix M' into the affine group is all we require.

Also noteworthy is the fact that there are infinitely many choices of similarity transformations S which will achieve the form given above: any invertible matrix with a constant row will work. This is not an issue for the theoretical results given in this paper, as they are independent of the particular choice of suitable similarity transformation S. However, in applications where practical considerations become important (such as finite sequence lengths) certain choices of S may outperform others. Some further comments on these matters are given in the discussion below.

Under an $m|n$ flattening we have $M_A = M_{s_1} \otimes M_{s_2} \otimes \ldots \otimes M_{s_m}$ as a Kronecker product of the m Markov matrices on the leaves A. After taking the similarity transformation $M_{s_j} \mapsto M'_{s_j} = S M_{s_j} S^{-1}$ on each factor, an elementary computation shows that (possibly after some simultaneous row and column permutations) we can extract an $(m(k{-}1){+}1) \times (m(k{-}1){+}1)$ sub-block which takes the form

$$\widehat{M}'_A := \begin{pmatrix} T_{i_1} & 0 & 0 & 0 & \cdots & v_{i_1} \\ 0 & T_{i_2} & 0 & 0 & \cdots & v_{i_2} \\ 0 & 0 & T_{i_3} & 0 & \cdots & v_{i_3} \\ \vdots & & & \ddots & & \vdots \\ 0 & 0 & 0 & \cdots & T_{i_m} & v_{i_m} \\ 0 & 0 & 0 & \cdots & 0 & 1 \end{pmatrix}. \tag{3-5}$$

Similarly, a corresponding expression exists for $\widehat{M}'_B$ as an $(n(k{-}1){+}1) \times (n(k{-}1){+}1)$ matrix.

As a simple example to convince the reader, consider the case of a two-fold Kronecker product:

$$M'_{12} \equiv M'_1 \otimes M'_2 = \begin{pmatrix} T_1 & v_1 \\ 0 & 1 \end{pmatrix} \otimes \begin{pmatrix} T_2 & v_2 \\ 0 & 1 \end{pmatrix} = \begin{pmatrix} T_1 \otimes T_2 & T_1 \otimes v_2 & v_1 \otimes T_2 & v_1 \otimes v_2 \\ 0 & T_1 & 0 & v_1 \\ 0 & 0 & T_2 & v_2 \\ 0 & 0 & 0 & 1 \end{pmatrix},$$

and define

$$\widehat{M}'_{12} := \begin{pmatrix} T_1 & 0 & v_1 \\ 0 & T_2 & v_2 \\ 0 & 0 & 1 \end{pmatrix}.$$

Applying additional Kronecker products requires some simultaneous row and column permutations to obtain the general form given in (3-5); however, the general construction and the reason for its existence should be clear.

In group theoretical language, we now observe (using elementary matrix multiplication) that (3-5) provides a *representation*[4] of the group $\times^m \mathrm{Aff}(k-1)$, where $\times^m \mathrm{Aff}(k-1)$ is realised by taking m direct products of $\mathrm{Aff}(k-1)$ via the Kronecker product $M'_A = M'_{s_1} \otimes M'_{s_2} \otimes \ldots \otimes M'_{s_m}$. Ultimately, the dimensional reduction we are about to describe hinges on the existence of this representation and is very much in accord with our general philosophy that emphasizes the importance of the transformation rules (3-2) and (3-3), as well as our previous explorations of Markov invariants [26, 20].

Theorem 1. *The form (3-5) provides a faithful ("one to one") representation of the direct product group* $\times^m \mathrm{Aff}(k-1)$ *with matrix entries that are linear in the matrix entries of the individual copies of* $\mathrm{Aff}(k-1)$.

Proof. The composition rule for two matrices in the affine group $\mathrm{Aff}(k-1)$ is revealed by the computation

$$\begin{pmatrix} T & v \\ 0 & 1 \end{pmatrix} \begin{pmatrix} U & w \\ 0 & 1 \end{pmatrix} = \begin{pmatrix} TU & Tw+v \\ 0 & 1 \end{pmatrix}.$$

In abstract terms, this can be understood by mapping each member $\left(\begin{smallmatrix} T & v \\ 0 & 1 \end{smallmatrix}\right) \in \mathrm{Aff}(k-1)$ to the pair (T, v) and defining the group product "$*$" via the rule[5]:

$$(T, v) * (U, w) = (TU, Tw + v). \tag{3-6}$$

In the case of the group $\times^m \mathrm{Aff}(k-1)$, each member may be represented as an m-tuple

$$((T_1, v_1), (T_2, v_2), \ldots, (T_m, v_m)),$$

and the product rule is provided by mimicking (3-6) in each entry of the tuple.

Correspondingly, the product of any two matrices of the form (3-5) is given by

$$\begin{pmatrix} T_{i_1} & 0 & 0 & \cdots & v_{i_1} \\ 0 & T_{i_2} & 0 & \cdots & v_{i_2} \\ \vdots & & \ddots & & \vdots \\ 0 & 0 & \cdots & T_{i_m} & v_{i_m} \\ 0 & 0 & \cdots & 0 & 1 \end{pmatrix} \begin{pmatrix} U_{i_1} & 0 & 0 & \cdots & w_{i_1} \\ 0 & U_{i_2} & 0 & \cdots & w_{i_2} \\ \vdots & & \ddots & & \vdots \\ 0 & 0 & \cdots & U_{i_m} & w_{i_m} \\ 0 & 0 & \cdots & 0 & 1 \end{pmatrix} = \begin{pmatrix} T_{i_1}U_{i_1} & 0 & 0 & \cdots & T_{i_1}w_{i_1} + v_{i_1} \\ 0 & T_{i_2}U_{i_2} & 0 & \cdots & T_{i_2}w_{i_2} + v_{i_2} \\ \vdots & & \ddots & & \vdots \\ 0 & 0 & \cdots & T_{i_m}U_{i_m} & T_{i_m}w_{i_m} + v_{i_m} \\ 0 & 0 & \cdots & 0 & 1 \end{pmatrix}.$$

Inspection of this product then shows that the matrices (3-5) form what is known as a representation of $\times^m \mathrm{Aff}(k-1)$, as required.

[4]The meaning of this will be given in the proof of Theorem 1.

[5]In fact, this the natural way to *define* $\mathrm{Aff}(k-1)$ in the first place.

Faithfulness of the representation — considered as an injective mapping from $\times^m \mathrm{Aff}(k-1)$ into the set of matrices of the form (3-5), is obvious.

Linearity of the matrix entries in (3-5) is also obvious.

□

Under the similarity transformation (3-4), the distribution P changes to (again expressed in tensor notation[6]):

$$P \to \underbrace{S \otimes S \otimes \ldots \otimes S}_{L=m+n \text{ products}} \cdot P. \tag{3-7}$$

Concurrently, the "transformed flattening" $\mathrm{Flat}'_{A|B}(P)$, is defined as

$$\begin{aligned}\mathrm{Flat}'_{A|B}(P) :&= \mathrm{Flat}_{A|B}(\underbrace{S \otimes S \otimes \ldots \otimes S}_{L=m+n \text{ products}} \cdot P) \\ &= \underbrace{(S \otimes S \otimes \ldots \otimes S)}_{m \text{ products}} \mathrm{Flat}_{A|B}(P) \underbrace{(S \otimes S \otimes \ldots \otimes S)^T}_{n \text{ products}}.\end{aligned}$$

Thus, the transformation rule (3-3) can then be expressed as

$$\mathrm{Flat}'_{A|B}(P) \to M'_A \mathrm{Flat}'_{A|B}(P) M'^T_B. \tag{3-8}$$

We now locate the subset of $(m(k-1)+1)$ rows and subset of $(n(k-1)+1)$ columns of $\mathrm{Flat}'_{A|B}(P)$ corresponding to the rows of $\widehat{M}'_A$ and the columns of $\widehat{M}'^T_B$ respectively (details are given in the appendix (A.3)), and define the "sub-flattening" $\widehat{\mathrm{Flat}}'_{A|B}(P)$ as a $(m(k-1)+1) \times (n(k-1)+1)$ sub-matrix of $\mathrm{Flat}'_{A|B}(P)$, and observe:

Corollary 1. *The sub-flattening $\widehat{\mathrm{Flat}}'_{A|B}(P)$ satisfies the transformation rule*

$$\widehat{Flat}'_{A|B}(P) \to \widehat{M}'_A \widehat{Flat}'_{A|B}(P)\ \widehat{M}'^T_B. \tag{3-9}$$

Proof. The result follows from (3-8) together with Theorem 1 and the observation that the form $\widehat{M}'_A$ (3-5) is block upper-triangular, and similarly the transpose $\widehat{M}'^T_B$ is block lower-triangular. □

Strikingly, according to Theorem 1, the sub-flattening transforms *bilinearly* in the parameters of the Markov matrices acting at the leaves of the phylogenetic tree (as opposed to the multi-linear rule exhibited in (3-2) and (3-3)). Additionally, we find that the contrasting rank conditions on $\mathrm{Flat}_{A|B}(P)$ discussed in the previous section are retained in a modified form on the sub-flattening $\widehat{\mathrm{Flat}}'_{A|B}(P)$. To state this result, we first need:

[6]Exactly how this affects the site pattern probabilities $p_{i_1 i_2 \ldots i_L}$ is given in the appendix (A.2).

Definition 1. We say that a phylogenetic model on a tree $\mathcal{T}$ is *generic* if the initial distribution $(\pi_i)_{i\in\kappa}$ at the root of the phylogenetic tree has no zero entries and the Markov matrix M_e on each edge $e \in \mathcal{T}$ has full rank.

It follows that:

Theorem 2. *The sub-flattening $\widehat{Flat}'_{A|B}(P)$ has rank at most $r(k-1)+1$, where $r \geq 1$ is the parsimony score for the split $A|B$ considered as a binary character at the leaves of $\mathcal{T}$. Up to the specified bound, maximal rank is attained in the generic case. In particular, in the generic case, the sub-flattening $\widehat{Flat}'_{A|B}(P)$ has rank k if the edge corresponding to the split $A|B$ occurs in $\mathcal{T}$ (since in this case $r=1$) and rank equal to or greater than $2(k-1)+1$ otherwise.*

Proof. The proof is given in Appendix A. □

These rank conditions are exactly analogous to the conditions for the flattenings given in [11] under the substitution $k^r \rightarrow r(k-1)+1$. We hence propose to use the sub-flattenings as an alternative, practical test for the existence of specific edges in a phylogenetic tree.

4. Discussion

Inspired by the previous approaches taken by [1, 6, 9, 11, 14] (as described at the end of the background section), we are currently exploring the computational means to exploit Theorem 2 in a practical setting. A potential obstruction is the change of basis required to convert an observed site pattern array F via the similarity transformation S. However, we note that this may be efficiently achieved by only computing the entries required for the sub-flattening and also by summing over only the non-zero entries of F. For large number of taxa L, F is an extremely large, and hence sparse, array. Without observing that the required transformation can be achieved efficiently there simply would be no way of computing the sub-flattening efficiently and hence the approach described in this paper would be of no practical use. From these observations, we claim this is not an insurmountable computational obstruction.

Additionally, in practical cases where the observed site pattern array F is obtained from *finite* sequence alignments, and hence is subject to standard statistical sampling error, it is quite possible that the specific choice of similarity transformation S (beyond having a constant row) will have an effect upon the performance of the method. Given this observation, and general best practice in numerical computation, it is likely that it is sensible to demand S to be an orthogonal matrix. At this stage we have not investigated this further (either theoretically or via simulation).

We suspect that the best way to test for low rank of the sub-flattenings will be via SVD, but at present it is unclear what the optimal numeric approach will be since the statistical properties under

multinomial sampling of the algebraic methods described here are unclear. Future work will explore these questions via simulations and testing on real data sets. Nonetheless, the results presented in this paper are expected to lead to a novel, computationally efficient phylogenetic method consistent with the general Markov model of molecular state evolution.

Acknowledgments. This work was inspired from a question Alexei Drummond put to Barbara Holland during her presentation at the New Zealand Phylogenetics Meeting, DOOM 2016. I would also like to thank the anonymous reviewer for their careful and substantive comments that lead to a greatly improved manuscript.

Funding. This work was supported by the Australian Research Council Discovery Early Career Fellowship DE130100423.

References

[1] Elizabeth S Allman, Laura S Kubatko, and John A Rhodes. Split scores: a tool to quantify phylogenetic signal in genome-scale data. *Systematic Biology*, page doi: 10.1093/sysbio/syw103, 2016. Online, 14 Nov 2017.

[2] Elizabeth S Allman and John A Rhodes. Phylogenetic ideals and varieties for the general Markov model. *Advances in Applied Mathematics*, 40(2):127–148, 2008.

[3] Andrew Baker. *Matrix groups: An introduction to Lie group theory*. Springer Science & Business Media, 2012.

[4] James D Bashford, Peter D Jarvis, Jeremy G Sumner, and Mike A Steel. U(1)× U(1)× U(1) symmetry of the Kimura 3ST model and phylogenetic branching processes. *Journal of Physics A: Mathematical and General*, 37(8):L81, 2004.

[5] David Bryant. Hadamard phylogenetic methods and the n-taxon process. *Bulletin of Mathematical Biology*, 71(2):339–351, 2009.

[6] Marta Casanellas and Jesús Fernández-Sánchez. Performance of a new invariants method on homogeneous and nonhomogeneous quartet trees. *Molecular Biology and Evolution*, 24(1):288–293, 2007.

[7] Marta Casanellas and Jesús Fernández-Sánchez. Relevant phylogenetic invariants of evolutionary models. *Journal de Mathématiques Pures et Appliquées*, 96(3):207–229, 2011.

[8] James A Cavender and Joseph Felsenstein. Invariants of phylogenies in a simple case with discrete states. *Journal of Classification*, 4(1):57–71, 1987.

[9] Julia Chifman and Laura Kubatko. Quartet inference from snp data under the coalescent model. *Bioinformatics*, 30(23):3317–3324, 2014.

[10] Jan Draisma and Jochen Kuttler. On the ideals of equivariant tree models. *Mathematische Annalen*, 344(3):619–644, 2009.

[11] Nicholas Eriksson. Tree construction using singular value decomposition. In Lior Pachter and Bernd Sturmfels, editors, *Algebraic Statistics for Computational Biology*, chapter 10, pages 347–358. Cambridge University Press, New York, 2005.

[12] Joseph Felsenstein. Evolutionary trees from DNA sequences: a maximum likelihood approach. *Journal of Molecular Evolution*, 17(6):368–376, 1981.

[13] Joseph Felsenstein. *Inferring phylogenies*, volume 2. Sinauer Associates Sunderland, 2004.

[14] Jesús Fernández-Sánchez and Marta Casanellas. Invariant versus classical quartet inference when evolution is heterogeneous across sites and lineages. *Systematic Biology*, 65(2):280–291, 2016.

[15] Walter M Fitch. Toward defining the course of evolution: minimum change for a specific tree topology. *Systematic Biology*, 20(4):406–416, 1971.

[16] Andrew R Francis. An algebraic view of bacterial genome evolution. *Journal of Mathematical Biology*, 69(6-7):1693–1718, 2014.

[17] Thomas R. Hagedorn. A combinatorial approach to determining phylogenetic invariants for the general model. Technical report, CRM-2671, 2000.

[18] Michael D Hendy, David Penny, and Mike A Steel. A discrete fourier analysis for evolutionary trees. *Proceedings of the National Academy of Sciences*, 91(8):3339–3343, 1994.

[19] Barbara R Holland, Peter D Jarvis, and Jeremy G Sumner. Low-parameter phylogenetic inference under the general Markov model. *Systematic Biology*, 62(1):78–92, 2013.

[20] Peter D Jarvis and Jeremy G Sumner. Adventures in invariant theory. *The ANZIAM Journal*, 56(02):105–115, 2014.

[21] Peter D Jarvis and Jeremy G Sumner. Matrix group structure and Markov invariants in the strand symmetric phylogenetic substitution model. *Journal of Mathematical Biology*, 73:259–282, 2016.

[22] Joseph E Johnson. Markov-type lie groups in $\mathrm{GL}(n, r)$. *Journal of Mathematical Physics*, 26(2):252–257, 1985.

[23] James A Lake. A rate-independent technique for analysis of nucleic acid sequences: evolutionary parsimony. *Molecular Biology and Evolution*, 4(2):167–191, 1987.

[24] Charles Semple and Mike Steel. *Phylogenetics*, volume 24. Oxford University Press, Oxford, 2003.

[25] Bernd Sturmfels and Seth Sullivant. Toric ideals of phylogenetic invariants. *Journal of Computational Biology*, 12(2):204–228, 2005.

[26] Jeremy G Sumner, Michael A Charleston, Lars S Jermiin, and Peter D Jarvis. Markov invariants, plethysms, and phylogenetics. *Journal of Theoretical Biology*, 253(3):601–615, 2008.

[27] Jeremy G Sumner, Jesús Fernández-Sánchez, and PD Jarvis. Lie Markov models. *Journal of Theoretical Biology*, 298:16–31, 2012.

[28] Jeremy G Sumner, Barbara R Holland, and Peter D Jarvis. The algebra of the general Markov model on phylogenetic trees and networks. *Bulletin of Mathematical Biology*, 74(4):858–880, 2012.

[29] Jeremy G Sumner and Peter D Jarvis. Entanglement invariants and phylogenetic branching. *Journal of Mathematical Biology*, 51(1):18–36, 2005.

[30] Jeremy G Sumner and Peter D Jarvis. Markov invariants and the isotropy subgroup of a quartet tree. *Journal of Theoretical Biology*, 258(2):302–310, 2009.

[31] Ziheng Yang. *Molecular evolution: a statistical approach*. Oxford University Press, Oxford, 2014.

Appendix A. Proof of Theorem 2

Our general approach to the proof will be to give conditions for when the rank of the sub-flattenings does or does not grow under phylogenetic divergence events. In particular, we will show that the rank of the sub-flattenings is unchanged after a phylogenetic divergence event which is "consistent" with the split under consideration (the precise meaning of this will become evident

below). Although related, our proof method is different in conception from the approach taken in [11] for obtaining the analogous conditions for the ranks of the full flattenings.

Definition 2. Given a rooted tree $\mathcal{T}$, consider the subtrees consisting of a vertex in $\mathcal{T}$ together with all of its descendants (including the case where the subtree consists of a leaf vertex only). Given a subset A of leaves, we say such a subtree *is* A*-consistent* if its leaves are a subset of A. We say an A-consistent subtree is *maximally* A*-consistent* if it is not itself a subtree of an A-consistent subtree. Similarly, given a split $A|B$ we say that a subtree is $A|B$*-consistent* if its leaves are a subset of A *or* B; together with the corresponding definition of *maximally* $A|B$*-consistent*.

An example is given in Figure 2.

FIGURE 2. A rooted tree with two maximally $A|B = \{2,3,4,6\}|\{1,5,7,8,9,10\}$ consistent subtrees indicated . The subtree with leaf set $\{3,4\}$ is A-consistent, but not maximally so.

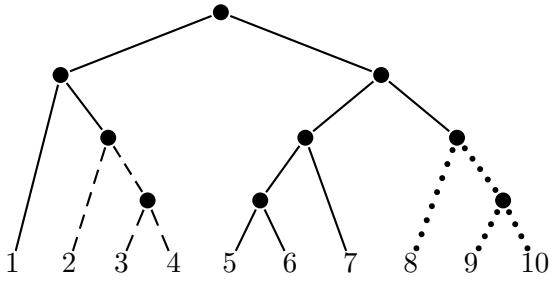

Lemma 1. *If P is a pattern distribution arising from a tree $\mathcal{T}$ under the general Markov model, the rank of the sub-flattening $\widehat{Flat}'_{A|B}(P)$ is independent of the size and/or structure of any $A|B$-consistent subtrees of $\mathcal{T}$.*

Proof. Consider the molecular state space $\kappa = \{1, 2, \ldots, k\}$ and a site pattern probability distribution $p_{i_1 i_2 \ldots i_L}$ on L taxa. Suppose this distribution arises under the general Markov model on the tree $\mathcal{T}$ and subsequently a time-instantaneous divergence event occurs causing, without loss of generality, a copy of the L^{th} taxon to be created. Under the usual assumptions of this model, this results in a new distribution $P^+ = (p^+_{i_1 i_2 \ldots i_L i_{L+1}})_{i_j \in \kappa}$ on an $L+1$ taxon tree $\mathcal{T}^+$, with

$$p^+_{i_1 i_2 \ldots i_L i_{L+1}} = \begin{cases} p_{i_1 i_2 \ldots i_L} \text{ if } i_L = i_{L+1}; \\ 0, \text{ otherwise.} \end{cases} \tag{A.1}$$

Consider a split $A|B$ and suppose taxon L is contained in B. Consider the new split $A|B'$ where the new taxon $L+1$ has been adjoined to B to produce $B' = B \cup \{L+1\}$. We will show that the sub-flattening $\widehat{\text{Flat}}'_{A|B'}(P^+)$ is obtained from $\widehat{\text{Flat}}'_{A|B}(P)$ by simply repeating $k-1$ columns.

Let S be any $k \times k$ matrix consistent with the similarity transformation (3-4). In particular, this means that the k^{th} row of S is constant, and, without loss of generality, we will assume this is a row of 1s, i.e. $S_{kj} = 1$ for $j = 1, 2, \ldots, k$. We denote the application of this similarity transformation to the site pattern distribution as

$$q_{i_1 i_2 \ldots i_L} := \sum_{j_1, j_2 \ldots, j_L \in \kappa} S_{i_1 j_1} S_{i_2 j_2} \ldots S_{i_L j_L} p_{j_1 j_2 \ldots j_L}. \tag{A.2}$$

We will refer to these quantities as the "q-coordinates".

Now suppose, without loss of generality, $A|B = \{1, 2, \ldots, m\}|\{m+1, m+2, \ldots, L\}$, and write $q_{i_1 i_2 \ldots i_m, j_1 j_2 \ldots j_n}$ to emphasize the flattening corresponding to this split. After locating the rows and columns which define the form (3-5), the $(m(k-1)+1) \times (n(k-1)+1)$ entries of the sub-flattening are seen to be given by

$$\widehat{\text{Flat}}'_{A|B}(P) = \left(\quad q_{i_1 i_2 \ldots i_m, j_1 j_2 \ldots j_n} \quad \right)_{\substack{i_a = k \text{ for all but at most one } a \in \{1,2,\ldots,m\}; \\ j_b = k \text{ for all but at most one } b \in \{1,2,\ldots,n\}.}} \tag{A.3}$$

We now consider the effect of the divergence rule (A.1) on the q-coordinates. Again we suppose that the divergence event occurs on the L^{th} taxon. As a consequence of (A.1), a short computation shows that

$$q^+_{i_1 i_2 \ldots i_m, j_1 j_2 \ldots j_{n-1} j_n j_{n+1}} = \sum_{j, j' \in \kappa} S_{j_n j} S_{j_{n+1} j} S^{-1}_{j j'} q_{i_1 i_2 \ldots i_m, j_1 j_2 \ldots j_{n-1} j'}.$$

To construct $\widehat{\text{Flat}}'_{A|B}(P^+)$ we must consider three cases (recalling that we are assuming $S_{kj} = 1$ for each $j = 1, 2, \ldots, k$):

(i.) Suppose $j_n = j_{n+1} = k$. Then

$$q^+_{i_1 i_2 \ldots i_m, j_1 j_2 \ldots j_{n-1} k k} = q_{i_1 i_2 \ldots i_m, j_1 j_2 \ldots j_{n-1} k}.$$

(ii.) Suppose $j_n = k$ and $j_{n+1} \neq k$. Then

$$q^+_{i_1 i_2 \ldots i_m, j_1 j_2 \ldots j_{n-1} k j_{n+1}} = q_{i_1 i_2 \ldots i_m, j_1 j_2 \ldots j_{n-1} j_{n+1}}.$$

(iii.) Suppose $j_n \neq k$ and $j_{n+1} = k$. Then

$$q^+_{i_1 i_2 \ldots i_m, j_1 j_2 \ldots j_{n-1} j_n k} = q_{i_1 i_2 \ldots i_m, j_1 j_2 \ldots j_{n-1} j_n}.$$

In particular, for each choice $j = 1, 2, \ldots k-1$,

$$q^+_{i_1 i_2 \ldots i_m, \underbrace{k k \ldots k k}_{n \text{ indices}} j} = q_{i_1 i_2 \ldots i_m, \underbrace{k k \ldots k}_{n-1 \text{ indices}} j} = q^+_{i_1 i_2 \ldots i_m, \underbrace{k k \ldots k}_{n-1 \text{ indices}} j k}.$$

Comparing to the general form (A.3), we see that the sub-flattening $\widehat{\text{Flat}}'_{A|B'}(P^+)$ is produced from the sub-flattening $\widehat{\text{Flat}}'_{A|B}(P)$ by simply repeating $k-1$ columns. This observation holds more

generally, independently of which taxon the divergence event occurs on. The only modification needed is when the divergence happens on the left side of the split $A|B$, in which case the new sub-flattening is obtained from the old by a repetition of rows rather than columns. Thus, if we place a new taxon into the same side of the split as the taxon it diverged from, the rank of the sub-flattening is preserved.

We now apply Corollary 1 to conclude that, in the generic case, further application of (full rank) Markov matrices at the leaves of the phylogenetic tree $\mathcal{T}^+$ also does not affect the rank of the sub-flattening.

These observations establish the lemma.

□

Now in order to determine the rank of an arbitrary sub-flattening, we may repeatedly apply Lemma 1 to reduce to the case where each $A|B$-consistent subtree is a single leaf. Assuming this situation, each leaf is then either (i) not part of a cherry, or (ii) part of a cherry where the two leaves in the cherry lie on complementary sides of the split $A|B$. A key feature of this situation is that we can label the descendants of every vertex (excluding the root) with complementary binary labels such that the leaf labels are consistent with the split $A|B$. For our purposes, we then consider this reduced case as arising from a sequence of divergence events from the base two-taxa case where, after each divergence event at a leaf, the two descendants are placed into complementary sides of the target split $A|B$. An example illustrating that this process is always possible given in Figure 3.

We use this process to establish:

Lemma 2. *Suppose $\mathcal{T}$ is a tree, suppose P is a distribution arising on $\mathcal{T}$ under the general Markov model, and suppose $A|B$ is a split such that the maximally $A|B$-consistent subtrees are all leaves. Then, in the generic case, the sub-flattening $\widehat{Flat}'_{A|B}(P)$ has maximal rank.*

Proof. Suppose such a reduced tree has q-coordinates $q_{i_1 i_2 \ldots i_m, j_1 j_2 \ldots j_n}$, and the n^{th} taxon in B diverges creating a new taxon which is adjoined to A to form the new split $A'|B$. Analogous to the previous situation, we have the new q-coordinates

$$q^+_{i_1 i_2 \ldots i_m i_{m+1}, j_1 j_2 \ldots j_n} = \sum_{j,j'=1}^{k} S_{i_{m+1} j} S_{j_n j} S^{-1}_{j j'} q_{i_1 i_2 \ldots i_m, j_1 j_2 \ldots j_{n-1} j'}.$$

From this we see that the additional $k-1$ rows in the sub-flattening $\widehat{\text{Flat}}'_{A'|B}(P^+)$ are obtained by setting $i_1 = i_2 = \ldots = i_m = k$, and taking $i_{m+1} = 1, 2, \ldots, k-1$ in

$$q^+_{kk \ldots k i_{m+1}, j_1 j_2 \ldots j_n} = \sum_{j,j'=1}^{k} S_{i_{m+1} j} S_{j_n j} S^{-1}_{j j'} q_{kk \ldots k, j_1 j_2 \ldots j_{n-1} j'},$$

FIGURE 3. Given a tree $\mathcal{T}$ and a split $A|B$ on its leaf set, leaves belonging to A are labelled by "+" and leaves in B are labeled by "−". The tree is reduced by removing any $A|B$-consistent subtrees, and binary labels are attached to the vertices (excluding the root) such that the descendants of each vertex obtain complementary labels and the leaf labels are consistent with the split $A|B$. In the case illustrated, the second step follows as a consequence of the two leaves that are not part of cherry.

Step 1. Reduce each maximally $A|B$ consistent subtree to a leaf.

Step 2. Label each internal vertex (excluding the root) consistently so descendants of internal vertices are distinctly labelled.

Step 3. Arbitrarily resolve any remaining ambiguities.

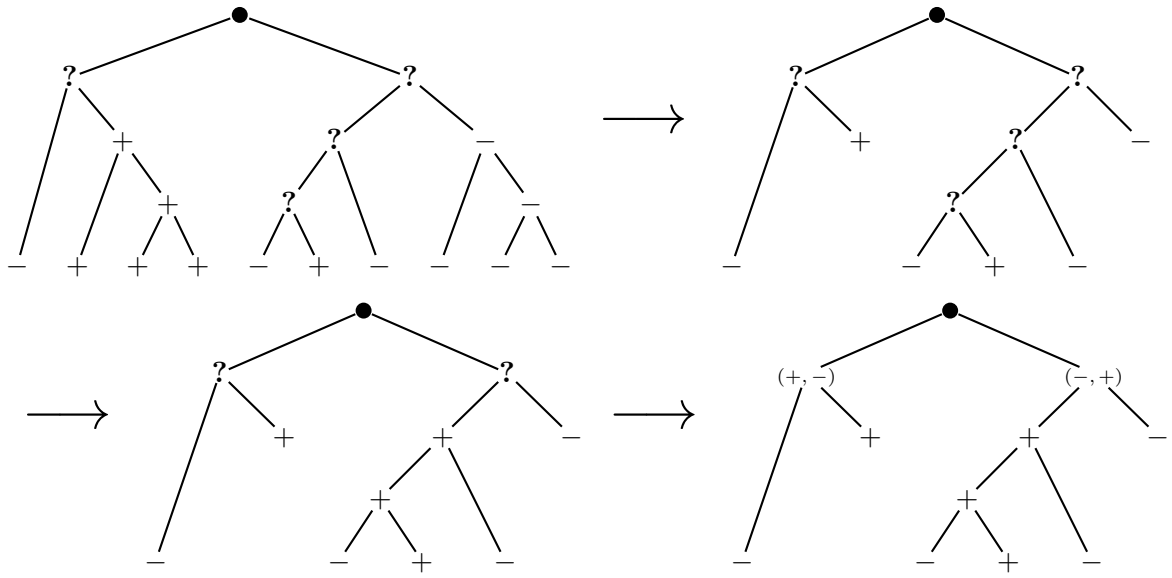

where the columns are indexed by choosing $b \in \{1, 2, \ldots, n\}$ so that at most a single $j_b \neq k$ at a time. In particular, if we choose $j_1 \neq k$ and $j_2 = j_3 = \ldots = j_n = k$ we have

$$q^{+}_{kk\ldots ki_{m+1},j_1kk\ldots k} = q_{kk\ldots k,j_1kk\ldots ki_{m+1}}.$$

Now for each choice $i_{m+1} = 1, 2, \ldots, k–1$ this expression gives q-coordinates which *do not* appear in the sub-flattening $\widehat{\text{Flat}}'_{A|B}(P)$ or any of the other rows of $\widehat{\text{Flat}}'_{A'|B}(P^+)$. It follows that any linear dependencies between the new and remaining rows in $\widehat{\text{Flat}}'_{A'|B}(P^+)$ would imply linear constraints on the q-coordinates on the original $m+n$ taxon tree. In turn, this would imply the existence of linear phylogenetic invariants for the general Markov model, which are known not to exist [17]. Therefore, the new rows appearing in $\widehat{\text{Flat}}'_{A'|B}(P^+)$ are linearly independent from the rest.

To complete the proof, we use induction on the base case of a two-taxon tree. To establish this base case, we show that, in the generic case, the two-taxon sub-flattening on the split $A|B = \{1\}|\{2\}$ has full rank $k = (k-1)+1$. This follows easily since, in the two-taxon case, the sub-flattening is equal to the transformed flattening, that is

$$\widehat{\mathrm{Flat}}'_{\{1\}|\{2\}}(P) = \mathrm{Flat}'_{\{1\}|\{2\}}(P) = S\mathrm{Flat}(P)_{\{1\}|\{2\}}S^{-1}.$$

Thus the sub-flattening is related by the similarity transformation S to the flattening $\mathrm{Flat}(P)_{\{1\}|\{2\}}$, which a standard argument shows can be expressed as

$$\mathrm{Flat}_{\{1\}|\{2\}}(P) = M_1 D(\pi) M_2^T,$$

where $D(\pi)$ is the diagonal matrix formed from the root distribution $\pi = (\pi_i)_{i\in\kappa}$. Clearly this matrix is full rank if M_1 and M_2 are full rank and π has no zero entries. Thus, in the generic case, the two-taxon sub-flattening $\widehat{\mathrm{Flat}}'_{1|2}(P)$ has full rank.

Induction on this base case establishes the lemma.

□

With these results in hand, Theorem 2 follows for arbitrary trees and splits by the following three steps:

(i) Apply Lemma 1 and clip off any $A|B$-consistent subtrees;

(ii) Apply Lemma 2; and

(iii) Use Fitch's algorithm [13, 15] to recognise that the minimum of the number of maximally A- and B-consistent subtrees is none other than the parsimony score for the split $A|B$ considered as a binary character at the leaves of $\mathcal{T}$.

J G Sumner, School of Physical Sciences, Mathematics, University of Tasmania, Private Bag 37, GPO, Hobart Tas 7001, Australia

E-mail address: Jeremy.Sumner@utas.edu.au